Trajectory Sonar Perception in the Ligurian Sea

Richard J. Rikoski[1], John J. Leonard[2], Paul M. Newman[3], and Henrik Schmidt[2]

[1] NSWC Panama City
Robotic Technologies Branch
richard.rikoski@navy.mil
[2] Massachusetts Institute of Technology
Department of Ocean Engineering
http://acoustics.mit.edu/GOATS
[3] University of Oxford
Robotics Research Group

Abstract. This paper describes a new technique for tracking locally curved unknown objects using sonar. The approach explicitly accounts for relevant robot dynamics. Objects are tracked by looking for temporal sequences of observations that fit a kinematic model. The method is illustrated using data from a synthetic aperture sonar gathered by a BPAUV at NATO SACLANT's GOATS 2002 experiment in the Ligurian Sea.

1 Introduction and State-of-Art

To navigate relative to features, avoid obstacles, recognize objects, and interact with or manipulate the underwater environment, substantial improvements in sonar perception need to be made.

Numerous groups have investigated sonar perception. Leonard and Durrant-Whyte [7] developed regions of constant depth (RCDs) to determine correspondence in a sonar scan. Nagatani *et al.* [10] developed the Arc Transversal Median method to find edges. Mataric and Brooks [2] developed a robot that maintained a fixed distance from walls by tracking without an explicit model. Kleeman and Kuc [5] used the method of images to distinguish corners and walls. Barshan and Kuc [1] developed a sonar that differentiated between convex corners and walls based on amplitude measurements. Kuc [6] developed a sonar that could recognize objects based on waveforms; the sonar adaptively positioned itself to search for unique aspects of the target. Peremans *et al.* [13] built a tri-aural sonar for estimating the range and bearing to a target. Wijk and Christensen [18,19,17] developed Triangulation Based Fusion, a method for mapping vertical edges using a sliding window of recent sonar measurements. Kleeman [4] designed a sonar that used double pulse coding to reject interference and perform classification.

Rikoski and Leonard developed a method for tracking locally curved objects in a two-dimensional environment [15]. This paper will extend that result to three dimensions.

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
2006		**00-00-2006 to 00-00-2006**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Trajectory Sonar Perception in the Ligurian Sea	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute of Technology,Department of Ocean Engineering,Cambridge,MA,02139**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
In Experimental Robotics IX, pages 557-570. 2006. U.S. Government or Federal Rights License

14. ABSTRACT
This paper describes a new technique for tracking locally curved unknown objects using sonar. The approach explicitly accounts for relevant robot dynamics. Objects are tracked by looking for temporal sequences of observations that fit a kinematic model. The method is illustrated using data from a synthetic aperture sonar gathered by a BPAUV at NATO SACLANT?s GOATS 2002 experiment in the Ligurian Sea.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Same as Report (SAR)**	**14**	

2 Trajectory Sonar Perception

Motivated by Marr's suggestion to "do what is possible, and proceed from there toward what is desirable" [9], the perception problem was broken down into a sequence of processing stages: signal processing (beamforming and matched filtering), detection, tracking, modeling, mapping, and object recognition. Each stage or competence had a well defined input, output, and transformation and was heavily grounded in physics [14]. The approach was designed to take advantage of delayed decision making [15,8]. Tracking, modeling, and object recognition were considered separate problems, as they happen on different time scales. This paper focuses on tracking.

We want to determine the minimum information necessary to determine correspondence between measurements of static features. We would like to do this without knowing what or where the feature is. This is appropriate because the modeling problem is greatly simplified if correspondence is known.

Trajectory Sonar Perception models geometric relationships between robots and features as potential fields. As a moving observer passes through a field, the observed rate of change of the field is known as the substantial derivative [11]. For a field Φ, the substantial derivative is defined as

$$\frac{D\Phi}{Dt} = \frac{\partial \Phi}{\partial t} + \dot{x} \cdot \frac{\partial \Phi}{\partial x}. \tag{1}$$

Given a prior measurement of the field Φ_{i-1}, Φ_i can be predicted using a Taylor series expansion of substantial derivatives

$$\Phi_i = \Phi_{i-1} + \frac{D\Phi_{i-1}}{Dt}\Delta t + \frac{D^2\Phi_{i-1}}{Dt^2}\frac{\Delta t^2}{2!} + \text{h.o.t.} \tag{2}$$

Assume a sensor that measures a geometric relationship between a robot and a target. Assuming that the geometric measurements vary continuously, and assuming the robot motion to be continuous, it is reasonable to apply the substantial derivative. A typical marine robot displaces 50-500kg, transmits its sonar 3-10 times per second, and is perturbed by waves with periods on the order of 5-30 seconds, making these assumptions reasonable.

Assume an initial measurement of a target $\begin{bmatrix} r & \theta & \phi \end{bmatrix}^T$. A subsequent measurement can be predicted through a Taylor series expansion of substantial derivatives.

$$\begin{bmatrix} r_1 \\ \theta_1 \\ \phi_1 \end{bmatrix} = \begin{bmatrix} r_0 + \frac{Dr}{Dt}\Delta t + \frac{D^2 r}{Dt^2}\frac{\Delta t^2}{2!} + \text{h.o.t.} \\ \theta_0 + \frac{D\theta}{Dt}\Delta t + \text{h.o.t.} \\ \phi_0 + \frac{D\phi}{Dt}\Delta t + \text{h.o.t.} \end{bmatrix} \tag{3}$$

The Taylor series for range includes a second order term. As will be shown later, the second substantial derivative of range is equivalent to the first substantial derivatives of azimuth and elevation. The second substantial derivative of range is

$$\frac{D^2r}{Dt^2} = \frac{\partial^2 r}{\partial t^2} + 2\dot{\mathbf{x}} \cdot \nabla \frac{\partial r}{\partial t} + \ddot{\mathbf{x}} \cdot \nabla r + \dot{\mathbf{x}}^T (\nabla \nabla^T r) \dot{\mathbf{x}}. \tag{4}$$

In the typical definition of the substantial derivative, $\dot{x}$ is the vector of vehicle velocities $\begin{bmatrix} u & v & w \end{bmatrix}^T$. For the purposes of this paper, $\dot{x}$ will be extended to include the roll, pitch, and yaw rates of the robot $\begin{bmatrix} P & Q & R \end{bmatrix}^T$. This is because the relative target bearing changes as the robot turns. Expanded, the substantial derivative of range becomes

$$\begin{aligned} \frac{Dr}{Dt} &= \frac{\partial r}{\partial t} + \frac{\partial r}{\partial x}\frac{\partial x}{\partial t} + \frac{\partial r}{\partial y}\frac{\partial y}{\partial t} + \frac{\partial r}{\partial z}\frac{\partial z}{\partial t} + \frac{\partial r}{\partial \xi_x}\frac{\partial \xi_x}{\partial t} + \frac{\partial r}{\partial \xi_y}\frac{\partial \xi_y}{\partial t} + \frac{\partial r}{\partial \xi_z}\frac{\partial \xi_z}{\partial t} \\ &= \frac{\partial r}{\partial t} + \frac{\partial r}{\partial x}u + \frac{\partial r}{\partial y}v + \frac{\partial r}{\partial z}w + \frac{\partial r}{\partial \xi_x}P + \frac{\partial r}{\partial \xi_y}Q + \frac{\partial r}{\partial \xi_z}R \end{aligned} \tag{5}$$

where ξ_x, ξ_y, and ξ_z denote rotation around the x, y, and z axes of the robot respectively.

By assuming the medium to be coherent [16], $\frac{\partial r}{\partial t}$, $\frac{\partial \theta}{\partial t}$,and $\frac{\partial \phi}{\partial t}$ can be neglected. This assumption implies that when the robot is stationary the measurements are stationary.

3 Universal Target Model

In order to calculate the substantial derivative of measurements, one must have a geometric model. In most prior work, point-like and planar objects have been treated as distinct. While maintaining discrete feature types is sound when there are relatively few, this approach will likely break down as more feature types are added and the differences between them is reduced. It would be desirable to have a single generic feature model that applies to all possible feature types.

By assuming targets to be locally curved, it is possible to accommodate both point-like and planar surfaces. If a locally curved surface has a center of curvature (x_c, y_c, z_c) and a radius of curvature ρ, then a point-like object has $\rho = 0$ and a planar surface has $\rho = \infty$. All other objects have finite non-zero radii of curvature. A convex object has $\rho > 0$, a concave object has $\rho < 0$. Putting the feature into robot coordinates (the robot is at $(0, 0, 0)$ facing in the x direction), but leaving the robot coordinates (x, y, z) in the equations to illustrate the derivatives, the predicted measurement is

$$\begin{bmatrix} r \\ \theta \\ \phi \end{bmatrix} = \begin{bmatrix} \sqrt{(x_c - x)^2 + (y_c - y)^2 + (z_c - z)^2} - \rho \\ \arctan(\frac{y_c - y}{x_c - x}) \\ \arccos(\frac{x_c - x}{\sqrt{(x_c - x)^2 + (y_c - y)^2 + (z_c - z)^2}}) \end{bmatrix} \tag{6}$$

Unfortunately, a measurement $\begin{bmatrix} r & \theta & \phi \end{bmatrix}^T$ cannot be directly inverted to yield the target model $\begin{bmatrix} x_c & y_c & z_c & \rho \end{bmatrix}^T$. The model has four degrees of freedom, but the

measurement has only three degrees of freedom. However, it is possible to calculate the substantial derivatives.

Consider the first term in the substantial derivative of range, $\frac{\partial r}{\partial x}u$. The partial derivative $\frac{\partial r}{\partial x}$ is

$$\frac{\partial r}{\partial x} = -\frac{x_c - x}{\sqrt{(x_c - x)^2 + (y_c - y)^2 + (z_c - z)^2}}. \tag{7}$$

Although the robot does not know (x_c, y_c, z_c), the robot can measure θ and ϕ. Substituting, it is evident that

$$\frac{\partial r}{\partial x} = -\cos(\theta)\sin(\phi). \tag{8}$$

By designing the sensor to measure θ and ϕ, it is possible to essentially measure the substantial derivative of range. The first substantial derivative of range of a locally curved object calculated using angular measurements is

$$\frac{Dr}{Dt} = -u\cos(\theta)\sin(\phi) - v\sin(\theta)\sin(\phi) - w\cos(\phi). \tag{9}$$

Notice that the first substantial derivative of range depends only on measurements and robot velocities (which can also be measured); it does not require globally referenced information or angular rates. Given velocity information, it can be calculated entirely from a previous measurement.

The derivatives of θ and ϕ, and the second substantial derivative of r are more difficult. The distance to the center of curvature $\sqrt{x_c^2 + x_c^2 + z_c^2}$ is squared in the denominator and cannot be removed entirely by substituting trigonometric identities. The range r cannot be substituted because $\sqrt{x_c^2 + x_c^2 + z_c^2}$ equals $r + \rho$, and ρ is not known. For instance:

$$\frac{D\theta}{Dt} = \frac{u\sin(\theta) - v\cos(\theta)}{\sin(\phi)\sqrt{(x_c - x)^2 + (y_c - y)^2 + (z_c - z)^2}} - \tag{10}$$

$$\frac{P\cos(\theta)\cos(\phi) + Q\sin(\theta)\cos(\phi)}{\sin(\phi)} - R.$$

The component with $\sqrt{(x_c - x)^2 + (y_c - y)^2 + (z_c - z)^2}$ in the denominator is the angular rate of the target due to robot translation. At broadside ($\theta = \frac{\pi}{2}$, $\phi = \frac{\pi}{2}$) the angular rate of the target is $\frac{u}{r_c}$, with r_c being the distance to the center of curvature. Intuitively, this makes sense because the angular rate of a satellite orbiting at a distance r_c from the earth's center and with velocity u would have the same angular rate. Unfortunately, the radius of curvature cannot be observed from a single measurement so this term cannot be directly predicted.

The second substantial derivative of range suffers from a similar problem. When a robot passes a target at broadside, the target appears to accelerate away from the robot at $\frac{u^2}{r_c}$. A satellite orbiting experiences a centripetal acceleration of $\frac{u^2}{r_c}$, so again

this is intuitively sensible. Of course, since the target is "center avoiding" this is, strictly speaking, a centrifugal effect.

What becomes evident is that by substituting the first substantial derivatives of the angles into the second substantial derivative of range, unwanted terms can be forced to drop out. In addition to removing the undesired r_c term, the angular rates can be forced to drop out. Looking at $\frac{D\theta}{Dt}$, it is evident that if the robot has a yaw rate R, the angular rate of the target will be $-R$. If the robot has velocity u and is passing the target at broadside on the right as it yaws, it will accelerate away from the target at uR. If the robot has a turning radius ρ_r, so that its yaw rate $R = \frac{u}{\rho_r}$, the robot accelerates away from the target at $\frac{u^2}{\rho_r}$, which is the robot's centripetal acceleration. Finishing the substitution yields the second substantial derivative in range

$$\frac{D^2r}{Dt^2} = -\dot{u}\cos(\theta)\sin(\phi) - \dot{v}\sin(\theta)\sin(\phi) - \dot{w}\cos(\phi) \tag{11}$$

$$+(u\sin(\theta)\sin(\phi) - v\cos(\theta)\sin(\phi))\frac{D\theta}{Dt}$$

$$-(u\cos(\theta)\cos(\phi) + v\sin(\theta)\cos(\phi) - w\cos(\phi))\frac{D\phi}{Dt}.$$

It is not possible to predict the second substantial derivative in range from one prior measurement. Given two prior measurements, however, it is possible to measure the angular substantial derivatives and then predict a third measurement.

4 Non-dimensional Analysis

It is necessary to determine when terms are observable for two reasons. First, it is important to determine how many terms should be kept in the Taylor expansion. Second, terms such as the angular rates must be observable if they are used to calculate other quantities such as the second substantial derivative of range.

The first question is whether the range changes enough ping to ping to justify using the first substantial derivative. It will be assumed that the sonar that has a ping rate f_s, a wavelength λ, and quarter wavelength resolution. In order for the first substantial derivative of range to be observable, the range must change more than a quarter wavelength between pings:

$$\frac{Dr}{Dt}\Delta t > \frac{\lambda}{4} \tag{12}$$

Since the substantial derivative of range is maximized when the target is in the robot's direction of travel, using V for the robot's speed, and substituting $\frac{1}{f_s}$ for Δt, this can be reduced to

$$\frac{V}{f_s} > \frac{\lambda}{4}. \tag{13}$$

This leads to the non-dimensional number

$$N_1 = \frac{V}{\lambda f_s}. \tag{14}$$

If N_1 exceeds $\frac{1}{4}$, then the range rate is observable. If the resolution were some other fraction of a wavelength, then N_1 would have to exceed that fraction. N_1 is also used to determine whether the heave and sway velocities need to be modeled. If, using the heave or sway velocity, the calculated N_1 value is less than $\frac{1}{4}$, those terms can be dropped from first substantial derivative of range. This is important if velocity information is not available, as it reduces the minimum number of features necessary for a velocity solution.

The second substantial derivative of range is observable if

$$\frac{\Delta t^2}{2!}\frac{D^2 r}{Dt^2} > \frac{\lambda}{4}. \tag{15}$$

It can be thought of as containing three components. The first component, described by N_2, analyzes the tendency of the range rate to change as a robot drives past a target(the previously mentioned centrifugal effect). The second component considers the contribution of the robot's angular rate to the second substantial derivative of range, and is described by N_3. The third component, the contribution of acceleration, has the non-dimensional number N_4. The non-dimensional numbers are

$$N_2 = \frac{V^2}{2f_s^2 \lambda r_{min}} \tag{16}$$

$$N_3 = \frac{V\Upsilon}{2f_s^2 \lambda} \tag{17}$$

$$N_4 = \frac{\dot{V}}{2f_s^2 \lambda}. \tag{18}$$

In N_3, Υ represents the maximum angular rate of the robot. If either N_2, N_3, or N_4 exceed the fraction describing the wavelength resolution of the sonar, the second substantial derivative of range is observable.

The angular rate of the robot is due to two effects. The target bearing can change because the robot passes the target, or because the robot turns. N_5 describes the contribution of robot translation, N_6 describes the contribution of turning.

$$N_5 = \frac{VL}{r_{min} \lambda f_s} \tag{19}$$

$$N_6 = \frac{\Upsilon L}{\lambda f_s} \tag{20}$$

In both N_5 and N_6, the aperture of the sonar L appears. It is assumed that the angular rate of the target is observable only if the change in the target bearing exceeds the resolution of the sonar. The broadside resolution of a line array, $\frac{\lambda}{4L}$ is used for

simplicity. If a robot has a sonar with an aperture L, the angular rate is observable if N_5 or N_6 exceed $\frac{1}{4}$.

The first six numbers describe robot effects in calm water. However, the ocean does not stand still. The effect of a current can be modeled through a straightforward application of N_1, but the effects of waves are slightly more complicated. Only linear waves will be considered in this paper. To avoid confusion with acoustic parameters, the parameters describing surface waves will use uppercase variables (because surface waves are "bigger" than acoustic waves). The relevant parameters are the angular frequency Ω, amplitude A, wave number K, and wavelength Λ.

In deep water, it is obvious that particles on the water surface move up and down as waves pass. Particles also have horizontal motions that are equal in in amplitude to the vertical motions, but 90° out of phase. Consequently, the particle motions are circular, and the paths are called particle orbits. The diameter of these orbits decay exponentially with depth. If a robot is neutrally buoyant and small with respect to the wavelength of the wave, it is reasonable to expect it to orbit like a particle. The non-dimensional number used to describe linear deep water waves is

$$N_7 = \frac{A\Omega e^{Kz}}{f_s \lambda}. \tag{21}$$

The z variable is referenced to the surface and underwater values are negative. At a depth of half a wavelength Λ, the amplitude of a particle orbit is reduced to $Ae^{-\pi}$. However, a moderate ocean wave can have a wavelength in excess of $100m$, so wave effects can extend considerably below the surface.

The shallow water case is more complicated. The vertical motion of water particles is restricted, as they cannot pass into the sea floor. Shallow water wave particles move in ellipses; at the bottom the particles only move horizontally. The non-dimensional number that describes the contribution of linear shallow water waves is

$$N_8 = \frac{gAK}{\Omega} \frac{\cosh(K(z+h))}{f_s \lambda \cosh(Kh)}. \tag{22}$$

For this number, h is the water depth and $g = 9.8 \frac{m}{s^2}$ (strictly speaking ocean waves are called gravity waves, because gravity is the restoring force).

5 Experimental Setup

In 2002, in conjunction with the NATO SACLANT research center in La Spezia, Italy, a set of missions were conducted as part of the GOATS 2002 experiment. The experiment was conducted using Caribou, an Odyssey III vehicle built by Bluefin Robotics. The Odyssey III is a streamlined vehicle with a vectored thruster. The vehicle featured a 16-element dual line array synthetic aperture sonar built by MIT doctoral candidates Don Eicksted and T. C. Liu, and post-doctoral fellow Wen Xu.

Since the drag on the cylindrical acoustic transmitter likely exceeded the vehicle drag by an order of magnitude, it was hidden inside the vehicle hull, which was designed to be acoustically transparent. The transmitter was aimed at broadside, and transmitted a $5 - 25$kHz linear chirp. The receiver, mounted on the nose, consisted of two parallel arrays of eight elements. The elements were spaced 10cm apart and the arrays were spaced 15cm apart.

Fig. 1. On the left, the robot, *Caribou*, being lowered into the water off the coast of Elba. On the right is a close up of the 16-element dual line array used to receive the sonar signals. The green spheres are individual hydrophones.

A characteristic along-track robot velocity was $1.5\frac{m}{s}$, a representative wavelength was 0.1m, and the ping rate $f_s = 3$Hz. Since $N_1 = 5$, the range rate was observable. Although direct sea state measurements were unavailable, wave induced heave and sway velocities between $0.1\frac{m}{s}$ and $0.25\frac{m}{s}$ were routinely observed. Since any velocity in excess of $0.075\frac{m}{s}$ would cause noticeable contribute to the range rate, the contributions of the induced heave and sway velocities were observable.

The robot typically operated 10m off the bottom. At that minimum range, $N_2 = 0.125$. The maximum yaw rate and acceleration were approximately $\Upsilon = \frac{10^\circ}{s} = 0.1745\frac{rad}{s}$ and $0.1\frac{m}{s^2}$ respectively. The other two measures of the second substantial derivative of range were $N_3 = 0.22$ and $N_4 = 0.06$. Although N_3 put the robot close to the transition, the second substantial derivative was neglected. Most targets were at least 20m away, and targets directly under the robot were usually masked by the powerful bottom reflection.

Since $N_5 = 0.35$ and $N_6 = 0.41$, the angular rates were barely observable. The beamformer used to produce the results in [12] was noisy enough to justify ignoring the angular rate. The beamformer used in this paper allowed the angular rate to occasionally be observable. Nevertheless, the information was not used. Consequently, the measurement prediction used in this paper was

$$\begin{bmatrix} r_1 \\ \theta_1 \\ \phi_1 \end{bmatrix} = \begin{bmatrix} r_0 - (u\cos(\theta)\sin(\phi) + v\sin(\theta)\sin(\phi) + w\cos(\phi))\Delta t \\ \theta_0 \\ \phi_0 \end{bmatrix} \tag{23}$$

The results presented here are from the *Caribou* mission initiated at 4:11pm on June 15, 2002, off the coast of Framura, Italy. In the GOATS data sets, this mission is referred to as SAS_15_6_2002____16_11. This is the same data set used in [12]. The mission lasted for roughly half an hour, and was conducted in approximately 20m of water about 500m from shore. The robot started on the surface, dove to roughly 10m, and then circled the targets. The vehicle path is shown in Fig 2.

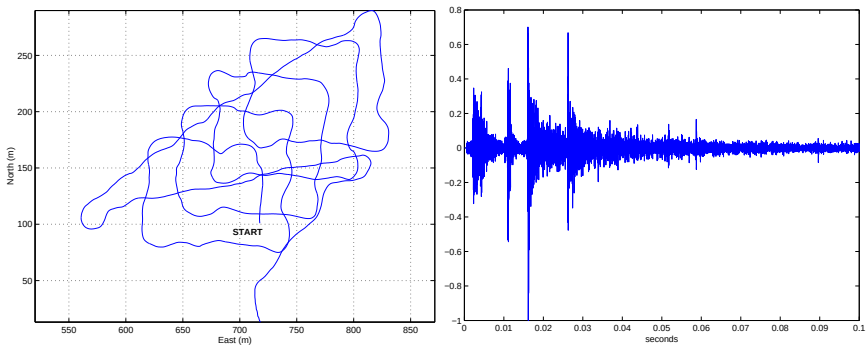

Fig. 2. The robot path is shown on the left. Since the sonar transmitted out the left side of the vehicle, the vehicle constantly turned left to keep the target field insonified. A typical matched filtered waveform is shown on the right. The pulse at the front of the waveform is a remnant of the outgoing signal. The next three pulses correspond to the reflection off the surface, the reflection off the bottom, and the multipath reflection off the surface and then the bottom. Potential targets are evident in the second half of the waveform.

6 Experimental Results

Starting with raw sonar signals, the robot systematically refined the data, building towards higher level constructs. The data was matched filtered, beamformed, and passed to a detector. The output of the detector was passed to feature tracker, which determined correspondence between pings of unknown objects using vehicle dynamics. The time series of raw detections and tracked features are shown in Fig 3 and Fig 4. Detections are shown in blue and measurements from tracked targets are shown in red.

In the time series of measurements, the detections at 3 meters are artifacts of the outgoing chirp. The two curves that are centered at approximately 10 meters are the reflections off the ocean surface and bottom. The two curves are mirror images of one another because, for a flat bottom, if the robot depth increases the robot altitude decreases. The measurements of the surface reflections appear to be noisier than measurements of the bottom because the ocean surface is not static.

The slowly varying line at 20m corresponds to the water depth. The signal reflects off the surface, off the bottom, and then back to the robot (or off the bottom, then

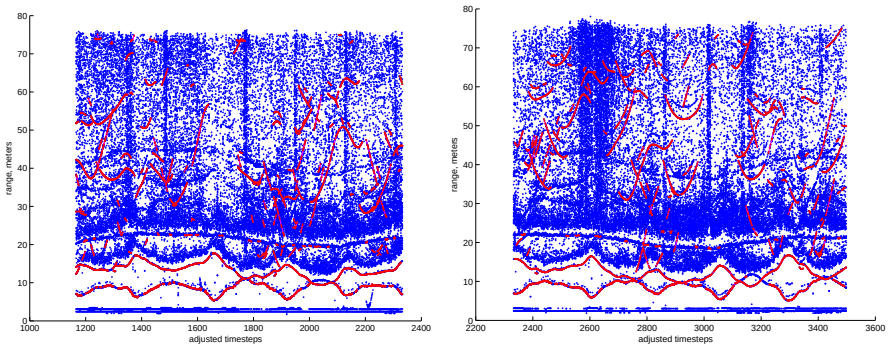

Fig. 3. This figure shows measurements plotted in measurement space. The x-axis is adjusted timestep, the y-axis is range in meters. Detections are plotted in blue, measurement trajectories are shown in red. The measurements on the right are for pings 1000 to 2000 (adjusted timesteps 1167 to 2334), the measurements on the right are for pings 2000 to 3000.

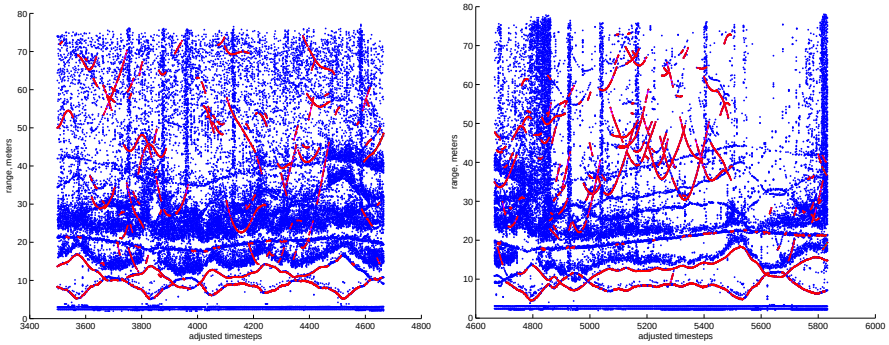

Fig. 4. On the left, the measurements and tracked features for pings 3000 to 4000. Pings 4000 to 5000 are shown on the right.

off the surface, then back to the robot, there are two paths). Centered around 30m are multi-path observations of the water surface and bottom. The surface multi-path reflection would bounce off the surface, off the bottom, off the surface, and then back to the robot. The reflections at approximately 40m are at twice the water depth, and correspond to two reflections off the surface and bottom. Each reflection off the surface introduces unmodeled noise into the sequence of measurements, since the water surface changes shape from ping to ping. Consequently, multi-path reflections are not tracked well.

The short vertical streaks are due to signals from the acoustic transponders used for navigation. The long wide streaks correspond to acoustic modem transmissions. Both signals tended to cause numerous false alarms at the level of the detector.

The detections and measurements from tracked features from pings 3000 to 4000 are projected into global coordinates in Fig 5. Since the altitude and depth reflections are from directly below and above the robot, they correspond to the vehicle's path.

Notice that, in general, this is an awkward representation for feature extraction. Also in Fig 5, there is a sequence of measurements corresponding to a mid-water target. This target was not tracked. This was likely due to one of two causes. First, it was inside the nearfield of the array. Bearing is a well defined construct for targets in the farfield of an array. However, when the farfield definition is applied to targets in the nearfield, the results range from imprecise to nonsensical.

Tracked targets start to appear after the bottom bounce, and extend out to 75m. Typical sequences of target measurements and predictions are shown in Figs 6 and 8. The robot is able to predict measurements to a fraction of wavelength. Differences between the predicted and actual measurements are primarily due to navigational errors. This is especially evident when comparing the predicted and actual range rates (Figs 7 and 9). The predicted range rate tends to lag the actual range rate. This is most likely because the state estimate from the navigation filter lags the true robot state.

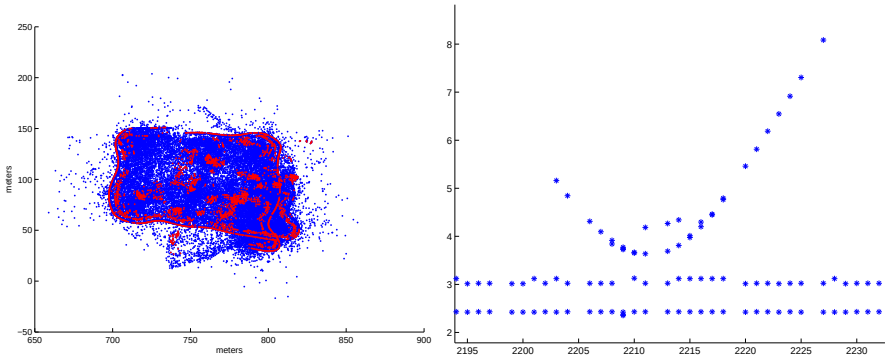

Fig. 5. On the right, the detections and tracked features from pings 3000 to 4000 are projected into cartesian coordinates. Clearly, this representation has less structure, making feature extraction more difficult. On the right is a feature that the algorithm failed to track. The curve on the right corresponds to a midwater object. Although the measurement are continuous, the measurements violate two assumptions. First, the second substantial derivative of range is non-negligible. Second, because the target is so close it's inside the nearfield of the array. The sequence of measurements at 2.5 and 3 meters are artifacts from the transmitted signal.

7 Conclusion

Using the constraint based on a universal feature model, a variety of features (point objects, concrete blocks, simple geometric shapes, the surface, the ocean floor) were tracked in a dynamic shallow water environment. The technique was able to find regularities in the data despite measurement and navigational uncertainty, and numerous interfering sources such as acoustic modems and long baseline navigation

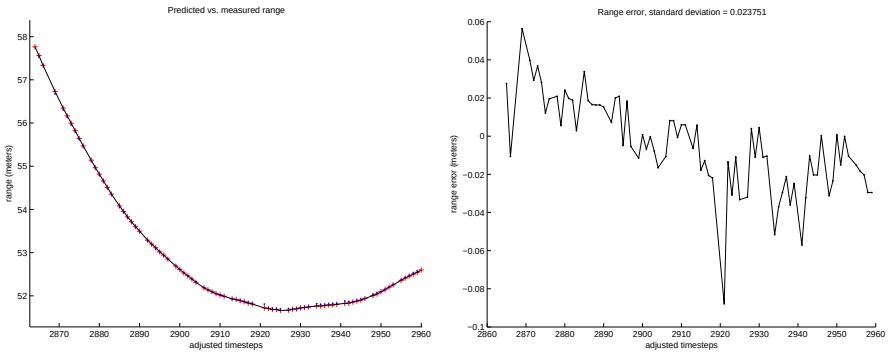

Fig. 6. On the right, a time series of measurements from a single target. The x-axis is adjusted timestep, the y axis is range. The measurements are plotted as red points connected by the black line, the predicted measurements are plotted in blue. The graph on the right shows the difference between the predicted and actual measurements. The errors are on the order of centimeters despite the robot moving at roughly $1.5\frac{m}{s}$ and the target being over 50m away.

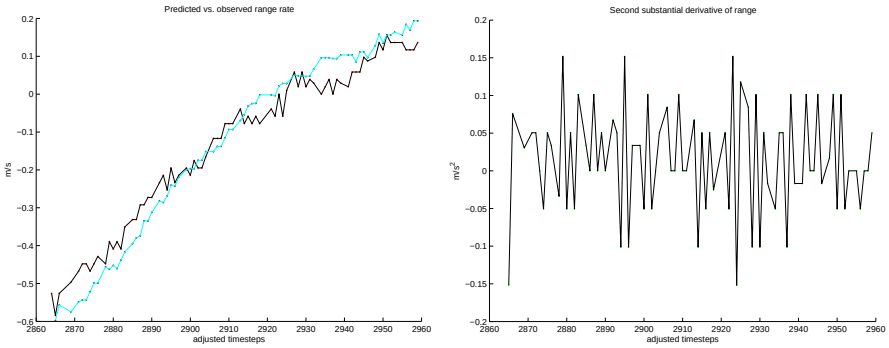

Fig. 7. The predicted and measured substantial derivatives of range are shown on the left. The range rate calculated from measurements is shown as red points connected by a black line, the range rate predicted using vehicle dynamics is shown as black points connected by a blue line. The second substantial derivative of range as calculated from sets of three measurements is shown on the right. Using non-dimensional analysis it was concluded that this derivative would not noticeably affect the time sequence of measurements from ping to ping. The data agrees with the non-dimensional analysis; the second substantial derivative is essentially zero mean noise.

beacons. The processed measurements were used by Newman as the input to a constant time SLAM algorithm [12].

There are four obvious research paths leading from this work. First, the work should be extended to dynamic objects. Second, techniques for creating geometric models of features given correspondence need to be investigated. Third, the approach needs to be extended beyond the Freedman model [3] to include rough surface

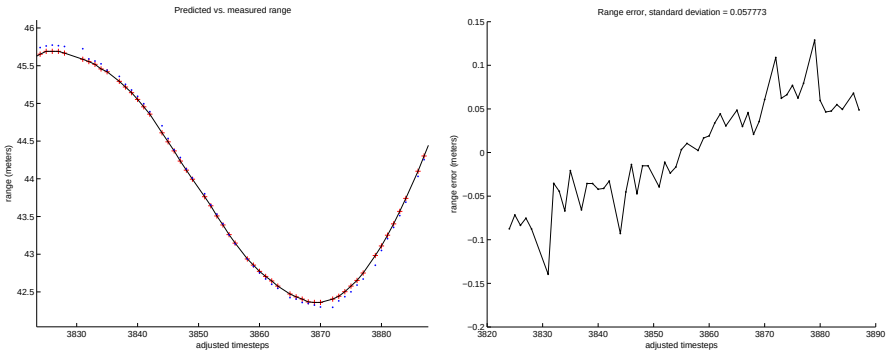

Fig. 8. A second time series of measurements from a target, and the difference between the predicted and observed range. In this case, the robot is turning while tracking the feature. This sequence of measurements has greater range errors, due to velocity errors as the robot turns.

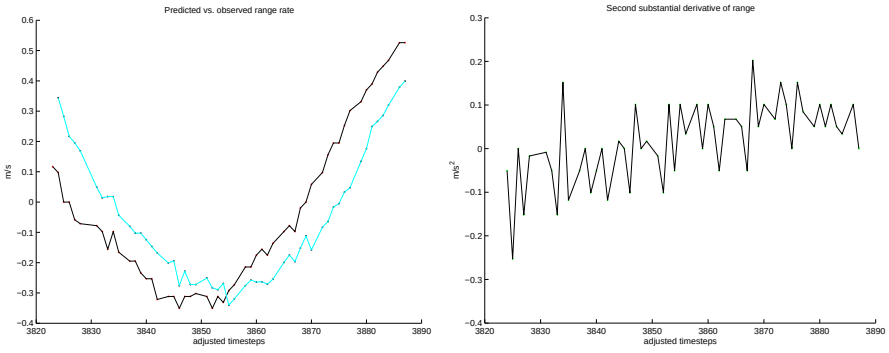

Fig. 9. Looking at the range rates, it is apparent why there were range errors in the previous figure. The navigation filter tends to lag the true state, that lag causes errors in the predicted range rate, which causes errors in the range prediction.

scattering. Finally, this analysis needs to be used to create a holistic design guide for marine robots using sonar.

Acknowledgement

This work was supported by the Office of Naval Research under grant N00014-97-1-0202.

We would like to thank the entire GOATS team. On the sonar side, Don Eickstedt, T. C. Liu, Wen Xu, and Joe Edwards designed and built the sensor used in this experiment, without them this work could never have happened. Rob Damus, Sam Desset and Jim Morash had the tough but vital job of fielding and maintaining the robots and their software. Matt Grund integrated the WHOI modems into the MIT vehicles, without which our capabilities would have been significantly reduced. The crew of the R.V. Alliance did a tremendous job of making the entire experiment

come off without a hitch, and keeping the scientific staff from hurting themselves. Finally, thanks to the Alliance's cook. The Alliance has a reputation for having the best food and wine cellar of any research vessel, but even so the cook far exceeded our expectations.

References

1. B. Barshan and R. Kuc. Differentiating sonar reflections from corners and planes by employing an intelligent sensor. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-12(6):560–569, June 1990.
2. R. A. Brooks. *Cambrian Intelligence: The Early History of the New AI*. The MIT Press, 1999.
3. A. Freedman. A mechanism of acoustic echo formation. *Acustica*, 12(1):10–21, 1962.
4. L. Kleeman. Advanced sonar with velocity compensation. *Int. J. of Robot. Res.*, 23(2):111–126, February 2004.
5. L. Kleeman and R. Kuc. Mobile robot sonar for target localization and classification. Technical Report ISL-9301, Intelligent Sensors Laboratory, Yale University, 1993.
6. R. Kuc. Fusing binaural sonar information for object recognition. In *IEEE/SICE/RSJ International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 727–735, 1996.
7. J. J. Leonard. *Directed Sonar Sensing for Mobile Robot Navigation*. PhD thesis, University of Oxford, 1990.
8. J. J. Leonard, R. J. Rikoski, P. M. Newman, and M. Bosse. Mapping partially observable features from multiple uncertain vantage points. *Int. J. of Robot. Res.*, 21(10/11):943–976, October 2002.
9. D. Marr. *Vision*. New York: W. H. Freeman and Co., 1982.
10. K. Nagatani, H. Choset, and N. Lazar. The arc-transversal median algorithm: an approach to increasing ultrasonic sensor accuracy. In *IEEE ICRA*, pages 644–651, 1999.
11. J. N. Newman. *Marine Hydrodynamics*. The MIT Press, 1977.
12. P. M. Newman, J. J. Leonard, and R. J. Rikoski. Towards constant-time SLAM on an autonomous underwater vehicle using synthetic aperture sonar. In *ISRR*, 2003.
13. H. Peremans, K. Audenaert, and C. J. M. Van. A high-resolution sensor based on tri-aural perception. *IEEE Transactions on Robotics And Automation*, 9(1):36–48, Feb. 1993. USA.
14. W. Richards, editor. *Natural Computation*. The MIT Press, 1988.
15. R. J. Rikoski. *Dynamic Sonar Perception*. PhD thesis, MIT, 2003.
16. S. J. Stanic, Ralph R. Goodman, Roger W. Meredith, Edgar Kennedy. Measurements of high-frequency shallow-water acoustic phase fluctuations. *IEEE J. of Oceanic Eng.*, 25(4):507–515, October 2000.
17. O. Wijk. *Triangulation Based Fusion of Sonar Data with Application in Mobile Robot Mapping and Localization*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2001.
18. O. Wijk, H. Christensen. Triangulation based fusion of sonar data with application in robot pose tracking. *IEEE Trans. Robotics and Automation*, 16(6):740–752, December 2000.
19. O. Wijk, H. I. Christensen. Localization and navigation of a mobile robot using natural point landmarks extracted from sonar data. *Robotics and Autonomous Systems*, 31:31–42, 2000.